Parametric and Probabilistic Model Checking of Confidentiality in Data Dispersal Algorithms (Extended Version)

Marco Baldi, Alessandro Cucchiarelli, Linda Senigagliesi, Luca Spalazzi, Francesco Spegni

Dipartimento di Ingegneria dell'Informazione
Università Politecnica delle Marche
Ancona, Italy

arXiv:1605.05930v2 [cs.CR] 29 Jun 2016

***Abstract*—Recent developments in cloud storage architectures have originated new models of online storage as cooperative storage systems and interconnected clouds. Such distributed environments involve many organizations, thus ensuring confidentiality becomes crucial: only legitimate clients should recover the information they distribute among storage nodes.**

In this work we present a unified framework for verifying confidentiality of dispersal algorithms against probabilistic models of intruders. Two models of intruders are given, corresponding to different types of attackers: one aiming at intercepting as many slices of information as possible, and the other aiming at attacking the storage providers in the network. Both try to recover the original information, given the intercepted slices.

By using probabilistic model checking, we can measure the degree of confidentiality of the system exploring exhaustively all possible behaviors. Our experiments suggest that dispersal algorithms ensure a high degree of confidentiality against the slice intruder, no matter the number of storage providers in the system. On the contrary, they show a low level of confidentiality against the provider intruder in networks with few storage providers (e.g. interconnected cloud storage solutions).

***Index Terms*—probabilistic model checking, parameters, security, confidentiality, information dispersal**

I. Motivation

Recently, the spreading of online storage services (such as iCloud, Dropbox, Skydrive, etc.) has seen a huge increase. According to the classical paradigm, the service provider buys or rents a large number of servers in which authorized clients are able to store their data. Recently, two other paradigms have emerged, viz. cooperative storage services (CSS) and federated (storage) clouds (FC).

In CSS, the storage capacity is provided directly by the clients themselves who, co-operating in the cloud, make their own storage facilities available to the others. This approach offers some evident advantages: first of all the service provider only needs a small number of servers, acting as coordinators for the access to the service. Secondly, increasing the number of users yields an increase of the storage capabilities. In this context, data is stored by the users, each acting as a *storage server*. In FC, organizations decide to agree on sharing their resources for solving common tasks. In this scenario, data is read from (and written to) several *storage providers*, each managing its set of storage servers behinds its cloud.

A fundamental requirement of CSS and FC is confidentiality: only the legitimate client should be able to recover the original information.

A consolidated solution to achieve confidentiality in such contexts is based on data dispersal. Dispersal algorithms provide a methodology for storing information in n distinct pieces, or slices, (dispersed) across multiple locations, so that redundancy protects the information in the event of a location outage, but unauthorized access at any single location does not provide usable information. Only the originator or a user which has access to, at least, k out of the n slices distributed among m available servers (or providers) can properly assemble and recover the complete information, without the need of any pre-shared encryption key. Instead, a client or attacker retrieving a number of slices lower than k is not able to get any information. This basic principle has been applied since the pioneer works by Shamir [1] and Rabin [2], and subsequently confirmed by McEliece and Sarwate [3] who disclosed the relationship with Reed-Solomon (RS) coding schemes.

Dispersal algorithms based on RS schemes have optimal performances, but they are subject to constraints due to the algebraic nature of the codes which practically limit the number of servers m, i.e., the level of dispersion. Since a large number of servers is convenient both from the client point of view (which can tolerate a greater number of failures) and from the security point of view (since the attacker needs to crack a larger number of nodes to steal the data), new solutions based on the Luby transform (LT) codes for the dispersal algorithm have been proposed [4]. These codes have no limit, in principle, on the value of n. On the other side, a characteristic feature of LT is that there are two thresholds k_1 and k_2 such that if an attacker has access to at least k_1 slices it has some probability of reconstructing the entire message, while if it has access to at least k_2 slices it has all the needed information to reconstruct the message. For this reason, one can see the RS coding schemes as special cases of LT.

Both coding schemes, RS and LT, can be combined with an all-or-nothing-transform (AONT) [5] realising AONT-RS and AONT-LT [6] data dispersal algorithms.

In this work we provide a unified formal framework to

model check the mentioned dispersal algorithms against different kinds of attackers trying to intercept slices and reconstruct the original message. We use the tool PRISM [7] to verify the degree of confidentiality of such algorithms. Since the problem is intrinsically parametric, we also want to identify suitable conditions under which the verification outcome holds for any number of storage providers in the cloud. To this aim, we repeatedly measured the probability of a confidentiality attack to understand how it varies w.r.t. other parameters, and in particular when the number of slices n increases.

Two different types of intruder are analyzed: the first one can eavesdrop the slices traveling to a server without interfering with the communication. The second type of intruder can violate some providers and retrieve all the slices they store. Both intruders are assumed to be *passive* and *probabilistic*, meaning that they can only read the exchanged information, and probabilities affect their capability of taking any action.

The organization of the paper is as follows: Section II defines the modeling languages and the models used for our analysis; Section III shows the formal analysis allowing us to measure the probability of an attack; Section IV comments the experimental data; Section V compares our work with the existing literature, while Section VI summarizes our results.

II. Modeling

Here, we formally describe a client process and two types of intruder. The main client responsibility is sending a sequence of *slices* to several distinct storage servers, following the specific dispersal algorithm. The set of slices constitutes the *message*. A *message* includes an actual content, viz. *message body*, and some extra information, viz. *message payload*, containing, among other things, the replicated information allowing to reconstruct the message body even in presence of faulty storage nodes. The main responsibility for both types of intruder is to intercept the traveling slices, and reconstruct the message body. The key difference between them is that the first type intercepts every slice independently from the previous ones. The second type of intruder, on the other side, attacks the storage providers and collect all the stored slices. The actions of the intruders are probabilistic.

The following parameters affect the system behavior: n is the number of slices that compose the message, m is the number of storage providers (or servers) in the system, c is the number of slices every server/provider can store, also called *capacity*, k_1 and k_2 are two thresholds such that if the attacker intercepts at least k_1 slices it has *some* probability of reconstructing the entire message body, while if it intercepts at least k_2 slices it has *all* the necessary information to reconstruct the message body. Two series of probabilities are used: a_i is the probability of intercepting a slice traveling to storage provider i, for the first attacker, and it is the probability of attacking the storage provider i, for the second attacker; x_j is the probability of reconstructing the entire message body, given j slices have been intercepted by the attacker. The relation among parameters are the following: $k_1 \leq k_2 \leq n$ and $n \leq m \cdot c$. Probabilities a_i are defined for $i \in [1, m]$, while probabilities x_j are defined for $j \in [k_1, n]$ and is such that: $\forall j \in [k_1, k_2 - 1].\ 0 < x_j < 1$ and $\forall j \in [k_2, n].\ x_j = 1$. Also: $\forall j \in [k_1, n].\ x_j \leq x_{j+1}$.

A. The system

Markov Decision Processes, or MDP for short, is a formalism allowing the definition of systems with probabilistic and non-deterministic actions. They are thus recognized as a good means to model randomized distributed systems: each process is described by its probabilistic transition function, and processes are interleaved by a non-deterministic scheduler . We briefly introduce MDPs using Baier and Katoen's notation [8].

Definition 1 (Markov Decision Process). *Assume a finite set of atomic propositions AP. A* Markov decision process *is a tuple $\mathcal{M} = (S, Act, Pr, \iota, L)$ where:*

- *$S = \{s_1, s_2, \ldots\}$ is a finite set of states,*
- *$Act = \{\alpha_1, \alpha_2, \ldots\}$ is a finite set of actions,*
- *$Pr : S \times Act \times S \to [0, 1]$ is a probabilistic transition function such that $\sum_{s' \in S} Pr(s, \alpha, s') \in \{0, 1\}$, for all $s \in S$ and $\alpha \in Act$;*
- *$\iota : S \to [0, 1]$ is the initial distribution probability of states, such that $\sum_{s \in S} \iota(s) = 1$;*
- *$L : S \to 2^{AP}$ is a labeling function.*

We write $(s, \alpha, p, s') \in Pr$ whenever $Pr(s, \alpha, s') = p$, for some suitable states s, s', action α and probability $p > 0$. We call *transition* any such tuple.

If the MDP is in state s, an action α is *enabled* if $(s, \alpha, p, s') \in Pr$, for some state s' and probability p. If an action α is enabled, than the probabilities among α-transitions must form a probability distribution: $\sum_{s' \in S} Pr(s, \alpha, s') = 1$. More than one action can be enabled in the same state s, thus the sum of probabilities of all transitions leaving state s sum up to the number of enabled actions. Let us remark that while probabilities in MDP could be real values, for algorithmic purposes in this work we constrain them to be rational values.

In the following we make use of the uniform probability distribution $\epsilon_m : [1, m] \to [0, 1]$ having the property: $\forall i \in [1, m].\ \epsilon_m(i) = \frac{1}{m}$. We write states of MDPs as configurations of some given set of variables V. Given a state s and a variable $v \in V$, we write $s.v$ to denote the value of the variable in that state. Given two states s and s' and a set of variables V, we write $s \equiv_V s'$ meaning that the values of variables in V are the same in both states: $\forall v \in V.\ s.v = s'.v$. Given a labeling $L : S \to 2^{AP}$, we say that L is *invariant* w.r.t. V iff $\forall s, s' \in S.\ s \equiv_V s' \Rightarrow L(s) = L(s')$.

Given MDPs $\mathcal{M}_1$ and $\mathcal{M}_2$, we will denote with $\mathcal{M}_1 \parallel \mathcal{M}_2$ the MDP resulting from their *synchronous composition*.

Let us list the graphical conventions used in this work to depict MDPs (see Fig. 1 and following):

- the circles represent the values of variables pc_c or pc_a;
- transitions have labels of the form: $\{p\}[\alpha]\gamma$ where $p \in [0, 1]$ is the probability, α is the action, and γ is a boolean formula mentioning the variables of the state.

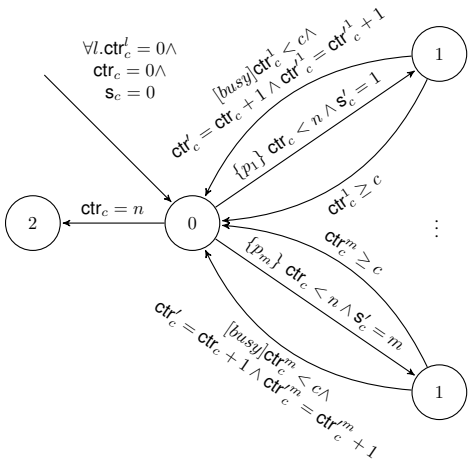
$\forall l.\mathsf{ctr}_c^l = 0 \wedge$
$\mathsf{ctr}_c = 0 \wedge$
$\mathsf{s}_c = 0$
$[busy]\mathsf{ctr}_c^l < c \wedge$
$\mathsf{ctr}'_c = \mathsf{ctr}_c + 1 \wedge \mathsf{ctr}'^l_c = \mathsf{ctr}^l_c + 1$
$\{p_1\}\ \mathsf{ctr}_c < n \wedge \mathsf{s}'_c = 1$
$\mathsf{ctr}_c = n$
$\mathsf{ctr}_c^l \geq c$
$\mathsf{ctr}_c^m \geq c$
$\{p_m\}\ \mathsf{ctr}_c < n \wedge \mathsf{s}'_c = m$
$\mathsf{ctr}'_c = \mathsf{ctr}_c + 1 \wedge \mathsf{ctr}'^m_c = \mathsf{ctr}^m_c + 1$
$[busy]\mathsf{ctr}_c^m < c \wedge$
2
0
1
1

Fig. 1. The MDP CLIENT

For the sake of brevity, we use notation $\exists i \in [a,b].\ \{p(i)\}[\alpha]\gamma(i)$ as transition label, to denote a group of $(b-a+1)$ similar transitions, each obtained by replacing i with one of the natural values in the interval $[a,b]$. Following the PRISM notation, the boolean formula on a transition can refer to variables in the source state by their name (e.g. ctr_c), and to variables in the target state by their primed name (e.g. ctr'_c).

B. The client

Let CLIENT be the MDP encoding the client process described earlier (see Fig. 1). It has the following local variables:

- pc_c : track the progress of the process,
- s_c : the identifier of the next recipient server/provider,
- ctr_c : the total number of sent slices,
- ctr_c^i : the number of slices sent to server/provider i.

The MDP has a first block of m transitions from $\mathsf{pc}_c = 0$ to $\mathsf{pc}_c = 1$ picking a storage server (or provider) to store the slice; each transition is subject to some probability p_i. Next, the client either sends the slice to the selected recipient, if the latter has not reached its capacity, or it tries again picking another one. A total slice counter and a server/provider slice counter are increased whenever the slice is sent. The loop terminates when all the slices are sent.

Let us remark that the sending transition is labeled with a special action $busy$. This is used when building the synchronous composition of CLIENT with the MDP modeling the intruder, to synchronize the action of sending by the client and the action of intercepting by the intruder.

C. The slice attacker

Let us name SLICEATT the MDP encoding the first type of intruder. The reason for its name is that it tries to intercept every slice, independently from the previously intercepted ones. The intruder is given in Fig. 2. It has two local variables, viz. pc_a and ctr_a. The former tracks the progress of the attack, while the latter counts the number of intercepted slices at any given moment.

The figure shows that the attack progresses linearly: it starts by intercepting the first k_1 slices, each with probability a_i given that the slice is sent to server/provider i. Having intercepted less than k_1 slices, there is no possibility to reconstruct the message body. After intercepting k_1 slices the next chain of states repeatedly alternate these steps: first it tries to reconstruct the message body with probability x_j, given $j = \mathsf{ctr}_a$; if it fails it tries to intercept a new slice. The state $\mathsf{pc}_a = \mathsf{done}$ denotes that the intruder reconstructed the message body.

D. The provider attacker

Let PROVIDERATT be the MDP encoding the second type of intruder. The reason for its name is that it tries to obtain the credentials of the storage provider, and later it will intercept all the slices traveling towards that provider.

The intruder is depicted in Fig. 3. It has two local variables, viz. pc_a and ctr_a. Similarly to the previous intruder, the former variable models the progress of the attack, while the latter counts the number of intercepted slices.

The intruder has an initial chain of m states where it tries to attack every provider tossing a coin with probability a_i; if the attack is successful it sets a flag att_a^i for provider i. State $\mathsf{pc}_a = m$ is reached when all attack attempts are decided (some succeeded and some failed). In a loop the attacker synchronizes with the $busy$ action from the client that is sending a message, to intercept every slice sent to an attacked provider. From state $(\mathsf{pc}_a = m, \mathsf{ctr}_a = j)$ there is a transition to some state with $\mathsf{pc} = \mathsf{done}$ labeled with probability x_j, meaning that it has probability x_j to reconstruct the content of the message, given j intercepted slices.

III. PARAMETRIC FORMAL VERIFICATION

In Section II we have seen that the problem at our hands is intrinsically parametric. The model checking problem requires its input MDP to be finite, thus we must fix the system parameters. On the other hand, this means that the outcome of our formal verification holds only for the specific configuration of the parameters themselves. One of the common desiderata when doing parametric formal verification, is to prove *universal properties*, i.e. we should check whether some property holds for any configuration of parameters. In this work we are able to measure the confidentiality of dispersal algorithms for any number of storage providers in presence of a slice attacker, while in the case of the provider attacker the degree of confidentiality depends on the actual number of storage providers in the network. Before showing the detailed formal analysis, we report the needed formal ingredients.

A. Preliminaries

PCTL* is a temporal logic for describing qualitative and quantitative aspects of probabilistic systems. The grammar of *PCTL** formulae is the following:

$$\begin{aligned} \Phi &::= \mathit{true} \mid p \mid \Phi \wedge \Phi \mid \neg\Phi \mid \mathbb{P}_J(\varphi) \\ \varphi &::= \Phi \mid \varphi \wedge \varphi \mid \neg\varphi \mid X\varphi \mid G\varphi \mid F\varphi \end{aligned}$$

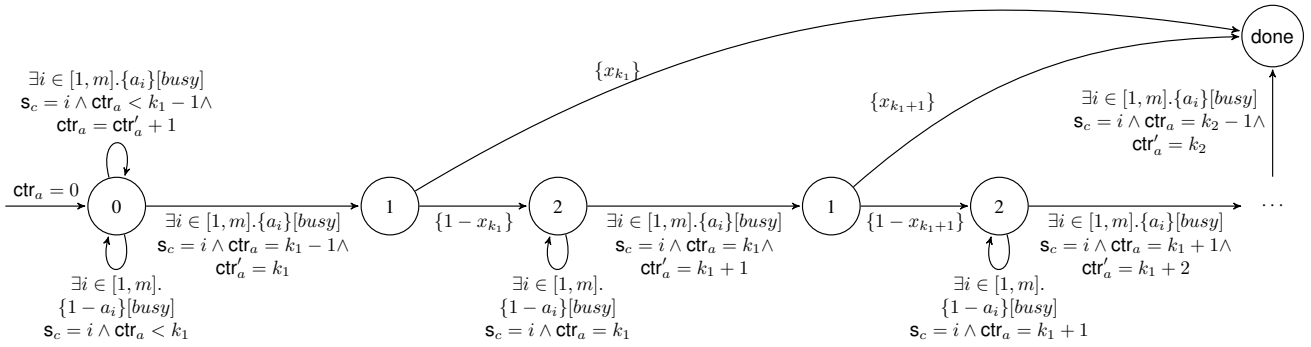
done
$\{x_{k_1}\}$
$\{x_{k_1+1}\}$
$\exists i \in [1,m].\{a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a < k_1 - 1 \wedge$
$\mathsf{ctr}_a = \mathsf{ctr}'_a + 1$
$\mathsf{ctr}_a = 0$
0
$\exists i \in [1,m].\{a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_1 - 1 \wedge$
$\mathsf{ctr}'_a = k_1$
1
$\{1 - x_{k_1}\}$
2
$\exists i \in [1,m].\{a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_1 \wedge$
$\mathsf{ctr}'_a = k_1 + 1$
1
$\{1 - x_{k_1+1}\}$
2
$\exists i \in [1,m].\{a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_2 - 1 \wedge$
$\mathsf{ctr}'_a = k_2$
$\exists i \in [1,m].\{a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_1 + 1 \wedge$
$\mathsf{ctr}'_a = k_1 + 2$
$\exists i \in [1,m].$
$\{1 - a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a < k_1$
$\exists i \in [1,m].$
$\{1 - a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_1$
$\exists i \in [1,m].$
$\{1 - a_i\}[busy]$
$\mathsf{s}_c = i \wedge \mathsf{ctr}_a = k_1 + 1$

Fig. 2. The MDP SLICEATT

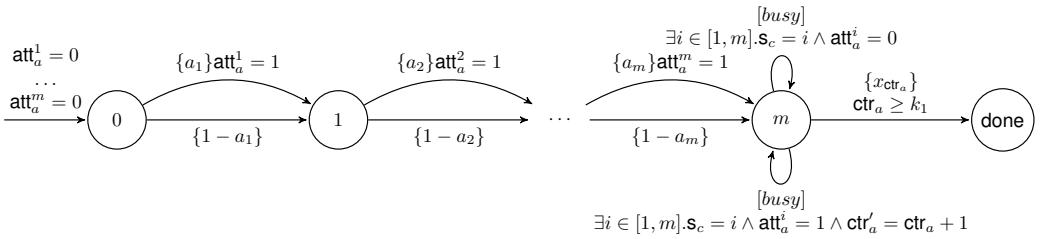

Fig. 3. The MDP PROVIDERATT

where $p \in AP$ and $J \subseteq [0,1]$ is a rational interval. Terms of Φ are *state formulae*, while terms of φ are *path formulae*.

A thorough description of the logic satisfiability relation is beyond the aims of this paper, since the subject is covered by several textbooks (e.g. see [8, Ch. 10.4]). Intuitively, formula $G\varphi$ holds w.r.t. some path iff every state visited in the path satisfies the sub-formula φ. Formula $F\varphi$, instead, holds w.r.t. some path iff some visited state satisfies sub-formula φ. The state formula $\mathbb{P}_{[a,b]}(\varphi)$ holds w.r.t. state s iff the sub-formula φ holds in all paths starting from s with some probability $p \in [a,b]$. Given an MDP $\mathcal{M}$, let us write $\mathcal{M} \models \Phi$ to express that *all the initial states* of $\mathcal{M}$ satisfy the property Φ.

Given a *PCTL** *path* formula φ and an MDP $\mathcal{M}$, there exist polynomial time algorithms computing the minimum and maximum probabilities of φ w.r.t. all the initial states of $\mathcal{M}$ [8, Ch. 10.6]. In the following we will write $\mathcal{P}_{min}(\varphi, \mathcal{M})$ and $\mathcal{P}_{max}(\varphi, \mathcal{M})$ to denote such computed probabilities. [1]

From the definitions of $\mathcal{M} \models \Phi$, $\mathcal{P}_{min}$ and $\mathcal{P}_{max}$, the following fact holds immediately.

Proposition 1 ([8, Ch. 10.6])**.** *Given any MDP $\mathcal{M}$, any PCTL* path formula φ, and any $0 \leq a \leq b \leq 1$, then:*

$$\mathcal{M} \models \mathbb{P}_{[a,b]}\varphi \iff a \leq \mathcal{P}_{min}(\varphi, \mathcal{M}) \wedge b \geq \mathcal{P}_{max}(\varphi, \mathcal{M})$$

[1]Note that, in general, given any *PCTL** formula φ and MDP $\mathcal{M}$, it is possible that $\mathcal{P}_{min}(\varphi, \mathcal{M}) \neq \mathcal{P}_{max}(\varphi, \mathcal{M})$. This is a consequence of the sequence of non-deterministic choices that can be taken in the executions of $\mathcal{M}$, each leading to a (possibly) different probability outcome associated to φ. This motivates the interest in discovering the minimum and maximum probabilities with which φ holds in $\mathcal{M}$.

Example 1. *Assume two MDPs $\mathcal{M}_1 :=$ CLIENT $\|$ SLICEATT and $\mathcal{M}_2 :=$ CLIENT $\|$ PROVIDERATT. Assume a proposition* HACKED $\in AP$ *and labeling L_1 of $\mathcal{M}_1$ (resp. L_2 of $\mathcal{M}_2$) such that* HACKED $\in L_1(s)$ *(resp.* HACKED $\in L_2(s)$*) iff $s.\mathsf{pc}_a = \mathsf{done}$. We can measure the likelihood of breaking the confidentiality requirement of $\mathcal{M}_1$ (resp. $\mathcal{M}_2$) computing $\mathcal{P}_{min}(F(\text{HACKED}), \mathcal{M}_1)$ and $\mathcal{P}_{max}(F(\text{HACKED}), \mathcal{M}_1)$ (resp. $\mathcal{P}_{min}(F(\text{HACKED}), \mathcal{M}_2)$ and $\mathcal{P}_{max}(F(\text{HACKED}), \mathcal{M}_2)$).*

Due to Proposition 1, the *probabilistic model checking* problem may assume two different flavours:

- *qualitative*: take as input an MDP $\mathcal{M}$ and a formula Φ, and return *true* iff $\mathcal{M} \models \Phi$;
- *quantitative*: take as input an MDP $\mathcal{M}$ and a path formula φ, and compute $\mathcal{P}_{min}$ and $\mathcal{P}_{max}$.

Here we use the *quantitative* probabilistic model checking.

Given an MDP $\mathcal{M}$, one can show that two states s and s' are indistinguishable, from a probabilistic point of view, if (i) every step taken from s is mimicked by some step taken from s', (ii) both steps end in equivalent states, and (iii) the viceversa is also true. This is captured by the notion of *probabilistic bisimulation*.

Definition 2 (**Probabilistic Bisimulation**, [8, Ex. 10.27])**.** *Given an MDP (S, Act, Pr, ι, L), a probabilistic bisimulation is an* equivalence relation $R \subseteq S \times S$ *such that, for any $s, s' \in S$, $R(s,s')$ iff:*

- $L(s) = L(s')$*, and*
- $Pr(s, \alpha, X) = Pr(s', \alpha, X)$*, $\forall \alpha \in Act, X \in {}^S/_R$*

where ${}^S/_R$ represents the quotient set *of S by R.*

Given two states s, t, let us write $s \approx_R t$ if $R(s,t)$ for some probabilistic bisimulation R. When R is clear from the context, we may omit it. The given definition of probabilistic bisimulation helps establishing that two states of the same MDP cannot be distinguished. It is possible to use probabilistic bisimulation to check whether different MDPs are indistinguishable. Given two MDPs, $\mathcal{M}_1 = (S_1, Act_1, Pr_1, \iota_1, L_1)$ and $\mathcal{M}_2 = (S_2, Act_2, Pr_2, \iota_2, L_2)$, let $S := S_1 \uplus S_2$ be the disjoint union of the state sets. Let us write $\mathcal{M}_1 \approx_R \mathcal{M}_2$ iff $AP_1 = AP_2$ and there exists a bisimulation relation $R \subseteq S \times S$ and $\iota_1(X) = \iota_2(X)$ for each $X \in {}^{S}/_{R}$. It is known that bisimilar MDPs satisfy the same *PCTL** formulae.

Theorem 1 ([8, Ex. 10.27])**.** *Given two MDPs $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1 \approx \mathcal{M}_2$, then $\mathcal{M}_1 \models \Phi$ iff $\mathcal{M}_2 \models \Phi$, for any $\Phi \in$ PCTL*.*

Proposition 1 and Theorem 1 yield the following.

Corollary 1. *Given two MDPs $\mathcal{M}_1$,$\mathcal{M}_2$ such that $\mathcal{M}_1 \approx \mathcal{M}_2$ and any PCTL* path formula φ, then $\mathcal{P}_{min}(\varphi, \mathcal{M}_1) = \mathcal{P}_{min}(\varphi, \mathcal{M}_2)$ and $\mathcal{P}_{max}(\varphi, \mathcal{M}_1) = \mathcal{P}_{max}(\varphi, \mathcal{M}_2)$.*

B. Parameter abstraction for the slice attacker

Let us consider the MDP $\mathcal{M}^{slice} = \text{CLIENT} \parallel \text{SLICEATT}$. In order to generalize our verification for any number of storage servers, we group the latter in *channels*, i.e. collections of servers that are indistinguishable.

Formally we define a channel as a triple (x, g, a) where $x \in \mathbb{N}_{>0}$ is the number of servers in the channel, $g : [1, x] \to [0, 1]$ is a probability distribution and $a \in [0, 1]$ is a probability value. Intuitively, x is the number of storage servers belonging to the channel, g is the probability distribution of picking any server in the channel when sending a slice, conditioned by the fact that the current channel has been chosen, and a is the probability of attacking any server belonging to the channel. It is easy to see that if a channel has size one, it can only be defined as follows: $(1, \epsilon_1, a)$, for some $a \in [0, 1]$.

In this analysis we assume that any storage server can host any number of slices. We also fix a subset of the model variables: $V = \{\mathsf{pc}_c, \mathsf{ctr}_c, \mathsf{pc}_a, \mathsf{ctr}_a\}$.

Given an Markov Decision Process $\mathcal{M}$, let us write $\mathcal{M}_f((x_1, g_1, a_1), \ldots, (x_k, g_k, a_k))$ denoting a copy of it where storage servers are grouped in the given k channels.

Our claim is that if we check formulae that look only at variables in V (and in particular that do not look at variables s_c and ctr^i_c, for $i \in [1, m]$), then a system with one server per channel is indistinguishable from a system with an arbitrary number of servers per channel. This means that the number of channels in the system defines an upper limit, or *cutoff*, to the size of the model to be verified.

Theorem 2. *Fix a positive number k and any probability distribution $f : [0, k] \to [0, 1]$. For all $n_1, \ldots, n_k \in \mathbb{N}_{>0}$, distribution probabilities $g_j : [0, n_j] \to [0, 1]$ ($j \in [1, k]$), probability values $a_1, \ldots, a_k \in [0, 1]$, let $\mathcal{M}_1 = \mathcal{M}^{slice}_f((1, \epsilon_1, a_1), \ldots, (1, \epsilon_1, a_k))$ and $\mathcal{M}_2 = \mathcal{M}^{slice}_f((n_1, g_1, a_1), \ldots, (n_k, g_k, a_k))$. Let L_1 and L_2 be the respective labeling functions, and assume they are invariant with respect to V. Then: $\mathcal{M}_1 \approx \mathcal{M}_2$.*

For the sake of readability, the detailed proofs can be found in the appendix of this paper.

Next corollary follows from Theorem 2 and Theorem 1.

Corollary 2. *Given $\mathcal{M}_1$ and $\mathcal{M}_2$ as of Theorem 2, then for all $\Phi \in$ PCTL*:*

$$\mathcal{M}_1 \models \Phi \iff \mathcal{M}_2 \models \Phi$$

C. Variable abstraction for the provider attacker

Here we focus on the scenario in which the intruder can break one or more providers, thus accessing all the slices they store. Two applicative examples of this scenario are the cooperative storage systems, where each node joining the network can receive some of the slices, and the interconnected cloud, where a client may decide to split the message among several cloud storage providers. In the latter case we ignore the fact the cloud storage provider may further distribute the received slices among its own servers, and treat each such provider as a single server.

The model checking problem: $\mathcal{P}_{min}(\varphi, \text{CLIENT} \parallel \text{PROVIDERATT})$, where φ encodes our confidentiality requirements, remains a problem parameterized by the actual numbers of storage providers (similarly for $\mathcal{P}_{max}$).

On the other side, one can easily see that if every provider can host any number of slices (i.e. $c \geq n$) we can produce an indistinguishable model that is much smaller w.r.t. the original one, by simply dropping the variables $W = \{\mathsf{ctr}^i_c : i \in [1, m]\}$ from the local state of CLIENT. This reduces the problem to a feasible one, allowing us to experimentally measure the degree of confidentiality of the considered dispersal algorithm w.r.t. the number of slices n and the number of providers m. Call V the set of remaining variables (i.e. $V \cap W = \emptyset$).

Name CLIENT$'$ a copy of CLIENT whose state does not contain variables in W. Intuitively, this means that CLIENT$'$ does not check whether a provider reached its capacity, but this is not a limitation since we assumed that every storage provider can host any number of slices.

Theorem 3. *Assume MDPs $\mathcal{M}_1 =$ CLIENT $\parallel$ PROVIDERATT and $\mathcal{M}_2 =$ CLIENT$'$ $\parallel$ PROVIDERATT. Let L_1 and L_2 be the respective labeling functions, and assume they are invariant with respect to V. Then: $\mathcal{M}_1 \approx \mathcal{M}_2$.*

Corollary 3. *Given $\mathcal{M}_1$ and $\mathcal{M}_2$ as of Theorem 3, then for all $\Phi \in$ PCTL*:*

$$\mathcal{M}_1 \models \Phi \iff \mathcal{M}_2 \models \Phi$$

IV. EXPERIMENTS

Here we show how to use model checking for measuring the likelihood of breaking the confidentiality requirement against

systems of growing sizes. In particular the parameter n is increased among runs. We remark that any cloud system is characterized by its own sets of parameters (e.g. the probabilities of attacking the used providers may be part of their SLA). In our experiments we choose parameters arbitrarily, mainly for showing the feasibility of the approach and underlining the weakness of data dispersal algorithms in some cloud environment , viz. interconnected clouds. The checked specifications are taken from Example 1.

Thanks to the parametric abstractions explained in Section III, the verification outcomes using the intruder SLICEATT hold no matter the number of storage servers/providers in the system. In the case of intruder PROVIDERATT the results depend on the number of storage providers in the network.

The experiments were run on a machine Xeon Quad Core 2.3 Ghz with 8 GB RAM and Linux 2.6.32 64 bit. The points in the graphs correspond to distinct instances of the model checking problem, all requiring from few seconds to 60 minutes to complete.

In the first set of experiments we compare two different types of dispersal algorithms, those based on RS transforms against those based on LT transforms. For this analysis we fixed the number of channels/providers to 3 and the attacking probabilities to $a_1 = 0.1$, $a_2 = 0.2$, and $a_3 = 0.3$. In the RS case we assumed $k_1 = k_2 = 0.7 \cdot n$, while in the LT case we assumed $k_1 = 0.6 \cdot n$ and $k_2 = 0.8 \cdot n$. In the case of the LT transforms we assumed that the sequence of probabilities x_j, for $j \in [k_1, n]$ is defined as follows: let $\delta = \frac{1}{k_2 - k_1 + 1}$, $x_j = (\min(j, k_2) - k_1 + 1) \cdot \delta$. Intuitively, probabilities x_j grow linearly in the interval $[k_1, k_2 - 1]$ and then stabilize at 1 for values greater than or equal to k_2 .

In Fig. 4 and 5 we depict how the probability of breaking the confidentiality requirement varies w.r.t. n. Let us remark that $\mathcal{P}_{max}$ and $\mathcal{P}_{min}$ versions of the formula coincide in every point of the series. We also notice that the algorithm (under the given parameters) shows a high degree of confidentiality against SLICEATT, while it is sensibly less confidential against PROVIDERATT. Moreover, the confidentiality in the latter case, after some initial change, stabilizes and does not depend on the actual number of exchanged slices. The latter observation is not surprising since the intruder mainly attacks providers, and few providers will receive many slices, thus giving the intruder a high probability of guessing the message body.

The second block of experimental data compares the effect of different attack probabilities, viz. $a_1 = 0.1$, $a_2 = 0.2$, and $a_3 = 0.3$ (see Low) vs. $a_1 = 0.3$, $a_2 = 0.4$, and $a_3 = 0.5$ (see High). For this analysis we used only LT transforms and we fixed the number of channels/providers to 3. As before, $k_1 = 0.6 \cdot n$ and $k_2 = 0.8 \cdot n$ and the series x_j grows linearly as for the previous experiment. Fig. 6 and 7 summarize the model checking outcomes. As would be expected, higher probabilities of intercepting slices give the data dispersal algorithm a very low level of confidentiality against intruder SLICEATT. Again, the low number of providers causes a low degree of confidentiality against intruder PROVIDERATT.

Finally, a third set of experiments compares the effects of different numbers of channels/providers in the system and the results are given in Fig. 8 and 9. In both cases a LT transform was used, with $k_1 = 0.6 \cdot n$ and $k_2 = 0.8 \cdot n$. In one case we assumed 5 channels/providers in the network (see m5) and in the other 10 channels/providers (see m10). The probabilities a_i, for $i \in [1, m]$, are distributed uniformly in the interval $[0, 0.25]$ in both cases. x_j is defined as for the previous experiment. Once again the degree of confidentiality against PROVIDERATT is considerably lower than that against SLICEATT. We underline that the experiment with 10 providers could be verified only for a small number of slices ($n \leq 150$), before running out of memory. Even though this fact represents a scalability issue of the presented methodology, the verification outcomes are still of practical interest since interconnected (or federated) cloud solutions usually employ a limited number of storage providers. Confidentiality of data dispersal with a small number of storage nodes appears to be too weak against attacks directed to the storage provider.

V. RELATED WORK

Formal verification of security requirements has a long history. In this area, model checking plays a predominant role [9], [10], [11], [12]. The traditional approaches consist in model checking the security requirements of a system opposed to an adversary able to intercept, remove, modify the original messages as well as to inject new messages. In this respect, the Dolev-Yao intruder model [13] is considered the most general model (the worst case) [14] as it assumes a non-deterministic attacker in full control of the communication channels.

Traditional model checking, though, is not suitable for verifying security of cloud systems: it can only verify whether a system can be attacked or not. We *assume*, instead, that every component of a cloud system can be attacked with some degree of probability, and are more interested in measuring the likelihood of such attacks. This motivated us to define custom probabilistic intruder models, in place of the Dolev-Yao intruder. To the best of our knowledge, few authors used probabilistic model checking for measuring security of systems [15], [16], [17], [18].

It is well known that model checking techniques must face the state-explosion problem, that easily makes the verification of real-world protocols and systems unfeasible. To overcome this limitation, one looks for *abstraction techniques* [19] that reduce the description of the system to a feasible state-space, still preserving the relevant properties. Special forms of abstractions are required when the system state-space depends on given parameters and one wants to check whether some property holds for all values of such parameters [20], [21].

With regards to probabilistic models, several approaches use abstraction techniques. Legay et al. [22] and Nouri et al. [23], for example, collect traces of real or simulated systems. Next they sample them in order to build an MDP that abstracts the original system. Such technique avoids to build a complete analysis of all the traces for large-scale (or even infinite) systems. The larger is the sample, on the other side, the higher is the probability that the verification outcomes are correct.

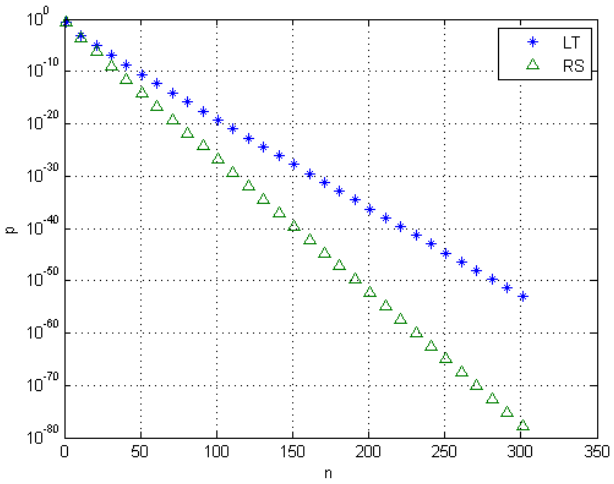

Fig. 4. Compare AONT-RS vs. AONT-LT using SLICEATT

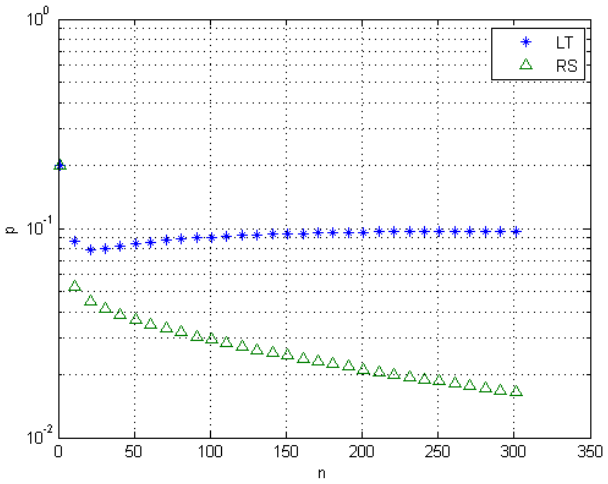

Fig. 5. Compare AONT-RS vs. AONT-LT using PROVIDERATT

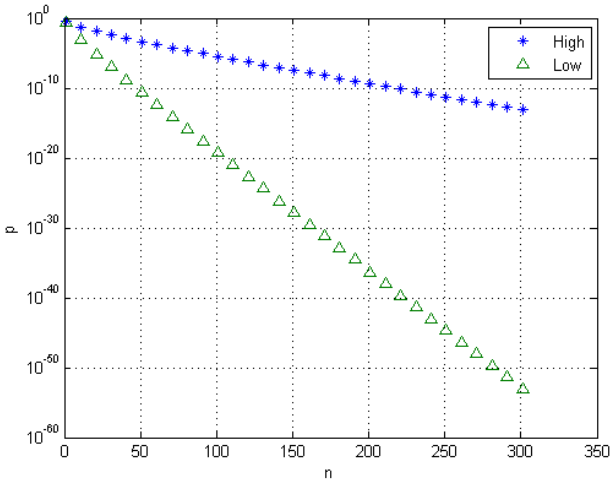

Fig. 6. Compare channel probability of attack using SLICEATT

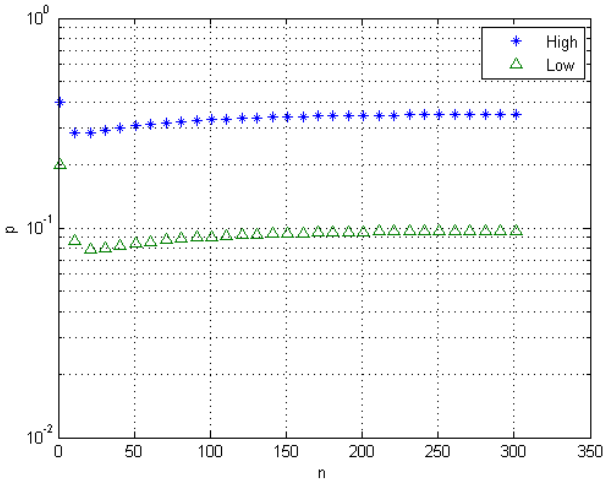

Fig. 7. Compare provider probability of attack using PROVIDERATT

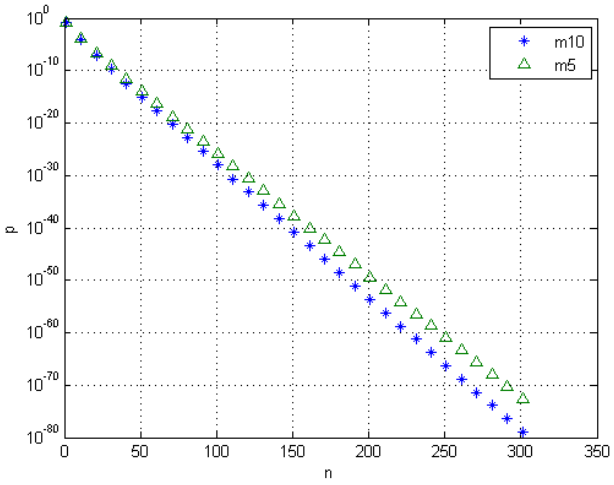

Fig. 8. Compare different number of channels using SLICEATT

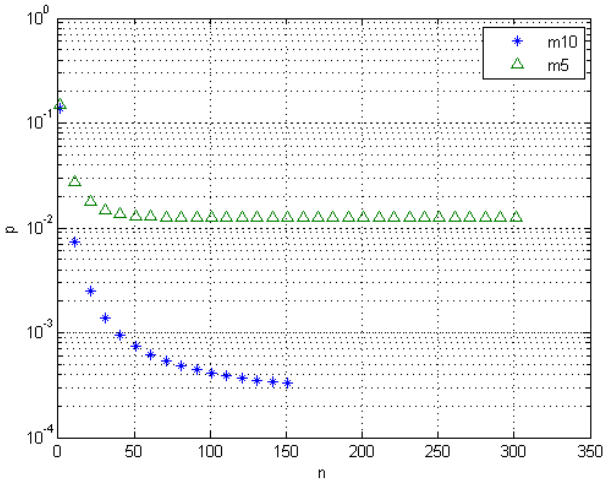

Fig. 9. Compare different number of providers using PROVIDERATT

Herd et al. [24], instead, proposed a trace sampling technique combined with trace fragmentation, i.e. only few fragments of a trace are considered. Abate et al. [25] proposed a method for transforming an MDP with an uncountable number of states into a Markov chain by means of a quotient-set based abstraction. The paper proves that the produced Markov chain approximates a probabilistic bisimulation of the original MDP.

Finally, let us remark that also the structure of the attacker may determine the feasibility of the verification of security properties. In our work we employed a passive intruder model, and indeed several authors agree that this is enough when analyzing confidentiality requirements. For example, Li and Pang [26], and Shmatikov [15] used passive intruders to verify anonymity of protocols, a special case of confidentiality. The latter work also considers probabilistic attacks. As far as we know, the use of a probabilistic passive attacker model for the analysis of data confidentiality is original.

VI. Conclusions

We presented a unified framework for the probabilistic model checking of a broad class of data dispersal algorithms in interconnected or cooperative cloud storage systems.

We verified confidentiality requirements of dispersal algorithms, checking the likelihood that an intruder has of intercepting slices of information and reconstruct the information.

In our framework we defined two types of probabilistic intruder, one tries to intercept each slice independently and the other attempts to attack the storage provider. In the former case the attack surface is the set of slices while in the latter it is the set of providers.

The problem is inherently parametric, since the CPU time and memory required to complete the verification are highly affected by several parameters, e.g. the number of slices used to split the information and the number of servers/providers.

By proving a probabilistic bisimulation property, we were able to generalize the results of the verification of confidentiality against the slice attacker to any number of servers in the network. The key observation, codified in our *channel abstraction*, is that any group of servers sharing the same probability of being eavesdropped may form a channel and behave like a single server hosting all the slices.

The analysis of the confidentiality against the provider attacker suggests, on the contrary, that classic data dispersal algorithms may not be the best solution to ensure confidentiality in interconnected cloud environments, unless the number of storage providers is considerably high.

We should remark that the conducted experiments fix some parameters to specific values. The conclusions thus are not fully generalizable w.r.t. such parameters. Among the modifiable parameters we remark that only k_1 and k_2 affect the state space and thus the complexity of the model checking problem.

We leave as future research the investigation of better abstractions suitable for the verification of the confidentiality of data dispersal algorithms against ProviderAtt on networks with many storage providers. That would improve the scalability of our framework to handle the case of cooperative storage systems.

References

[1] A. Shamir, "How to share a secret," *ACM Communications*, vol. 22, no. 11, 1979.

[2] M. O. Rabin, "Efficient dispersal of information for security, load balancing, and fault tolerance," *Journal of the Association for Computing Machinery*, vol. 36, no. 2, 1989.

[3] R. McEliece and D. Sarwate, "On sharing secrets and Reed-Solomon code," *ACM Communications*, vol. 24, 1981.

[4] M. Luby, "LT codes," in *Foundation of Computer Science 2002, Proceedings*, 2002.

[5] R. Rivest, "All-or-nothing encryption and the package transform," in *Fast Software Encryption 1997, Proceedings*, ser. LNCS, vol. 1267. Springer, 1997.

[6] M. Baldi, N. Maturo, E. Montali, and F. Chiaraluce, "AONT-LT: a data protection scheme for cloud and cooperative storage systems," in *HPCS 2014, Proceedings*, 2014.

[7] M. Kwiatkowska, G. Norman, and D. Parker, *PRISM 4.0: Verification of Probabilistic Real-Time Systems*. Springer, 2011.

[8] C. Baier and J.-P. Katoen, *Principles Of Model Checking*, 2008.

[9] D. Basin, C. Cremers, and C. Meadows, "Model checking security protocols," *Handbook of Model Checking*, 2011.

[10] C. Meadows, "Formal methods for cryptographic protocol analysis: Emerging issues and trends," *Journal on Selected Areas in Communications*, vol. 21, no. 1, 2006.

[11] F. Pagliarecci, L. Spalazzi, and F. Spegni, "Model checking grid security," *Future Gener. Comput. Syst.*, vol. 29, no. 3, 2013.

[12] M. Panti, L. Spalazzi, S. Tacconi, and S. Valenti, "Automatic verification of security in payment protocols for electronic commerce." in *ICEIS*, 2002.

[13] D. Dolev and A. C. Yao, "On the security of public key protocols," *Information Theory, IEEE Transactions on*, vol. 29, no. 2, 1983.

[14] I. Cervesato, "The Dolev-Yao intruder is the most powerful attacker," in *16th Annual Symp. on Logic in Computer Science (LICS)*, vol. 1, 2001.

[15] V. Shmatikov, "Probabilistic analysis of an anonymity system," *Journal of Computer Security*, vol. 12, no. 3, 4, 2004.

[16] G. Lenzini, S. Mauw, and S. Ouchani, "Security analysis of socio-technical physical systems," *Computers & Electrical Eng.*, vol. 47, 2015.

[17] S. Ouchani and M. Debbabi, "Specification, verification, and quantification of security in model-based systems," *Computing*, vol. 97, no. 7, 2015.

[18] F. Yang, G. Yang, and Y. Hao, "The modeling library of eavesdropping methods in quantum cryptography protocols by model checking," *International Journal of Theoretical Physics*, 2016.

[19] E. M. Clarke, O. Grumberg, and D. E. Long, "Model checking and abstraction," *ACM Trans. Program. Lang. Syst.*, vol. 16, no. 5, 1994.

[20] B. Aminof, T. Kotek, S. Rubin, F. Spegni, and H. Veith, "Parameterized model checking of rendezvous systems," in *CONCUR 2014, Proceedings*, vol. 8704. Springer, 2014.

[21] L. Spalazzi and F. Spegni, "Parameterized model-checking of timed systems with conjunctive guards," in *Verified Software: Theories, Tools and Experiments*. Springer, 2014.

[22] A. Legay, B. Delahaye, and S. Bensalem, "Statistical model checking: An overview," in *Runtime Verification*. Springer, 2010.

[23] A. Nouri, B. Raman, M. Bozga, A. Legay, and S. Bensalem, "Faster statistical model checking by means of abstraction and learning," in *Runtime Verification*. Springer, 2014.

[24] B. Herd, S. Miles, P. McBurney, and M. Luck, "Quantitative analysis of multiagent systems through statistical model checking," in *Engineering Multi-Agent Systems*. Springer, 2015.

[25] A. Abate, M. Kwiatkowska, G. Norman, and D. Parker, "Probabilistic model checking of labelled Markov processes via finite approximate bisimulations," in *Horizons of the Mind.* Springer, 2014.

[26] Y. Li and J. Pang, "Formalizing provable anonymity in Isabelle/HOL," *Formal Aspects of Computing*, vol. 27, no. 2, 2015.

APPENDIX

Corollary 1. *Given two MDPs $\mathcal{M}_1$,$\mathcal{M}_2$ such that $\mathcal{M}_1 \approx \mathcal{M}_2$ and any PCTL* path formula φ, then $\mathcal{P}_{min}(\varphi, \mathcal{M}_1) = \mathcal{P}_{min}(\varphi, \mathcal{M}_2)$ and $\mathcal{P}_{max}(\varphi, \mathcal{M}_1) = \mathcal{P}_{max}(\varphi, \mathcal{M}_2)$.*

Proof: Assume MDPs $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1 \approx \mathcal{M}_2$, and any path formula $\varphi \in$ *PCTL**. Let us name $a = \mathcal{P}_{min}(\varphi, \mathcal{M}_1)$ and $a' = \mathcal{P}_{min}(\varphi, \mathcal{M}_2)$. Assume, by contradiction and w.l.o.g., that $a > a' \geq 0$.

It is immediate that some $a \leq b \leq 1$ must exist such that: $\mathcal{M}_1 \models \mathbb{P}_{[a,b]}(\varphi)$, thus $\mathcal{M}_2 \models \mathbb{P}_{[a,b]}(\varphi)$ (by Theorem 1).

As we observed earlier (see Proposition 4), by definition of the model checking problem *PCTL**, $\mathcal{P}_{min}$ and $\mathcal{P}_{max}$ (see e.g. [8]), the following holds: $\mathcal{M} \models \mathbb{P}_{[x,y]}(\varphi)$ iff $x \leq \mathcal{P}_{min}(\varphi, \mathcal{M})$ and $y \geq \mathcal{P}_{max}(\varphi, \mathcal{M})$, for any MDP $\mathcal{M}$, *PCTL** path formula φ and rationals $0 \leq x, y \leq 1$.

In our case, and from $\mathcal{M}_2 \models \mathbb{P}_{[a,b]}(\varphi)$, it follows that:

$$a \leq \mathcal{P}_{min}(\varphi, \mathcal{M}_2) = a'$$

which in turn contradicts our assumption $a > a'$. This proves that $\mathcal{P}_{min}(\varphi, \mathcal{M}_1) = \mathcal{P}_{min}(\varphi, \mathcal{M}_2)$.

Using a symmetric argument, it is straightforward to prove that, given the corollary assumptions, $\mathcal{P}_{max}(\varphi, \mathcal{M}_1) = \mathcal{P}_{max}(\varphi, \mathcal{M}_2)$.

■

Theorem 2. *Fix a positive number k and any probability distribution $f : [0, k] \to [0, 1]$. For all $n_1, \ldots, n_k \in \mathbb{N}_{>0}$, distribution probabilities $g_j : [0, n_j] \to [0, 1]$ ($j \in [1, k]$), probability values $a_1, \ldots, a_k \in [0, 1]$, let $\mathcal{M}_1 = \mathcal{M}_f^{slice}((1, \epsilon_1, a_1), \ldots, (1, \epsilon_1, a_k))$ and $\mathcal{M}_2 = \mathcal{M}_f^{slice}((n_1, g_1, a_1), \ldots, (n_k, g_k, a_k))$. Let L_1 and L_2 be the respective labeling functions, and assume they are invariant with respect to V. Then: $\mathcal{M}_1 \approx \mathcal{M}_2$.*

Proof: Let $n := \sum_{i=1}^{k} n_i$. Let $H : [1, n] \to [1, k]$ be a surjective mapping. Intuitively, $H(j) = i$ iff j belongs to the i-th channel. If for some j, k, $H(j) = H(k)$ it means that $(p_j, q_j) = (p_k, q_k)$, i.e. they belong to the same channel, and we will write $j =_H k$. In this proof we call *transition template* a PRISM transition, possibly with indexed variables. Let us call Θ the set of the above transition templates. Given any $\theta \in \Theta$ and a state s, we write $\theta(s)$ to denote the actual MDP transition obtained by instantiating the variables in θ with their values in state s. It is evident that for any θ and s there exists a unique target state s'. For example: let θ be an indexed PRISM transition $[\alpha]\mathsf{foo}_i = 1 \to p : \mathsf{foo}'_i = 2$, and an MDP state s such that $s.\mathsf{foo}_1 = 1$ and $s.\mathsf{foo}_2 = 1$, then $\theta(s)$ is the MDP transition: (s, p, α, s') such that $s'.\mathsf{foo}_1 = 2$, and $s'.\mathsf{foo}_2 = 2$.

Given a transition $\tau = (s, \alpha, p, s') \in Pr$ we will denote with *pre*(τ) the source state s and with *post*(τ) the target state s'.

Name S_1 the set of states of $\mathcal{M}_1$, S_2 the set of states of $\mathcal{M}_2$, and S their disjoint union. Let us define the relation $R \subseteq S \times S$ as follows: $R(s_1, s_2)$ iff the following properties hold:

- $s_1 \equiv_V s_2$
- $s_1.\mathsf{s}_c =_H s_2.\mathsf{s}_c$
- $\forall i \in [1, m]$. $s_1.\mathsf{ctr}_i^1 = \sum_{j=1}^{n_i} s_2.\mathsf{ctr}_i^j$

In order to show that R is a probabilistic bisimulation, let us first observe that R is an equivalence relation (i.e. it is transitive, reflexive and symmetric).

Secondly, let us underline that $L(s_1) = L(s_2)$ follows by our definition of *AP* and the fact that $\forall v \in V$. $s_1.v = s_2.v$.

Finally, let us show that: $Pr(s_1, \alpha, X) = Pr(s_2, \alpha, X)$, for all $X \in S/R$ and all $\alpha \in$ *Act*.

Let us consider two cases: either (C1) $X = [s_1]_R$, or (C2) $X \neq [s_1]_R$. Notice that by definition of X and R, $[s_1]_R = [s_2]_R$.

Case (C1) is possible iff there is a transition τ satisfying all the following conditions:

- *pre*(τ) $\equiv_V$ *post*(τ) i.e. τ does not change any variable $v \in V$ (otherwise $s_1, s_2 \notin X$);
- $\forall i \in [1, n]$. *pre*(τ) $\equiv_{\{\mathsf{ctr}_i\}}$ *post*(τ), i.e. τ does not change any counter (otherwise the counters can only increase, thus their sum increase, and thus $s_1, s_2 \notin X$);
- either *pre*(τ) = *post*(τ) or *pre*(τ) $\not\equiv_{\{\mathsf{s}_c\}}$ *post*(τ) i.e. it is either a self-loop or it changes variable s_c.

Inspecting all the transitions in the model, there is none that satisfies the above conditions, thus this case is *impossible*.

Let us split case (C2) in subcases. Either (C2.1) X differs from $[s_1]$ only for variables in V, or (C2.2) X differs from $[s_1]$ for variables in $V \cup \{\mathsf{ctr}_i, \mathsf{s}_c\}$.

Case (C2.1) is possible for the following transition templates:

- any transition of SLICEATT, or
- the CLIENT transition template: $(\mathsf{pc}_c = 0, \mathsf{ctr}_c = \mathsf{n}) \to (\mathsf{pc}'_c = 3, \mathsf{s}_c = 0)$

Let us observe the following facts, for any $\theta \in \Theta$:

- $\exists q_1 \in X, p \in [0, 1]$. $\theta(s_1) = (s_1, p, \alpha, q_1) \Rightarrow \exists q_2 \in X$. $\theta(s_2) = (s_2, p, \alpha, q_2)$;
- $\exists q_2 \in X, p \in [0, 1]$. $\theta(s_2) = (s_2, p, \alpha, q_2) \Rightarrow \exists q_1 \in X$. $\theta(s_1) = (s_1, p, \alpha, q_1)$;

This means, intuitively, that whenever τ is applicable to state s_1 and reaching some state $q_1 \in X$, it is also applicable to the equivalent state (by R) s_2 and reaches some state q_2 equivalent to q_1. This can be proven checking all transition templates in Θ: their final statuses q_1, q_2 have the same values for variables in V and they didn't change the values of variables $\{\mathsf{s}_c, \mathsf{ctr}_i\}$ w.r.t. s_1, s_2. Since the correspondence preserves the value of probability p, it follows that:

$$\begin{aligned} Pr(s_1, \alpha, X) &= \sum_{\substack{q \in X, \theta \in \Theta, \\ \theta(s_1) = (s_1, p, q)}} p \\ &= \sum_{\substack{q' \in X, \theta' \in \Theta, \\ \theta'(s_2) = (s_2, p, q')}} p \quad = Pr(s_2, \alpha, X) \end{aligned}$$

Case (C2.2): X differs from $[s_1]$ for variables in V and for variables in $\{\mathsf{s}_c, \mathsf{ctr}_c^i\}$. The possible transition templates in this case are all CLIENT templates:

- $[\mathsf{busy}](\mathsf{pc}_c = 1, \mathsf{s}_c = i, \mathsf{ctr}_c^i < c) \rightarrow (\mathsf{pc}'_c = 0, \mathsf{ctr}'_c = \mathsf{ctr}_c + 1, \mathsf{s}'_c = 0)$
- $(\mathsf{pc}_c = 0, \mathsf{ctr}_c < n) \rightarrow p_i : (\mathsf{pc}'_c = 1, \mathsf{s}_c = i)$

In the case of the first template again we can directly check that:

- $\exists q_1 \in X, p \in [0,1].\ \theta(s_1) = (s_1, p, \alpha, q_1) \Rightarrow \exists q_2 \in X.\ \theta(s_2) = (s_2, p, \alpha, q_2)$;
- $\exists q_2 \in X, p \in [0,1].\ \theta(s_2) = (s_2, p, \alpha, q_2) \Rightarrow \exists q_1 \in X.\ \theta(s_1) = (s_1, p, \alpha, q_1)$;

Take all transitions induced by the first template: the final statuses satisfy the property $q_1 \equiv_W s_1$ and $q_2 \equiv_W s_2$, where $W = V \cup \{\mathsf{s}_c\}$. When variable ctr_c^j changes, for some j, it increases by one. In this case the same template θ can be applied to state s_2 and variable ctr_c^h can increase its value by one (with same probability) for some other channel h, such that $h =_H j$.

In the case θ is the second template, let us call $n = \sum_{i=1}^{k} n_i$, i.e. n represents the total number of servers belonging to some channel in the system. Let us define the set $A_s = \{\tau : \tau \in \theta, \exists q.\ Pr(s, \alpha, q) > 0\}$, i.e. A_s contains all the transitions that are enabled in s and are induced by template θ. Let us define the sets: $A_s^i = \{\tau : \tau \in \theta, \tau = (s, p, \alpha, q), p > 0, \langle \mathsf{ctr}_c^j < \mathsf{c} \rangle \in \mathit{guard}(\tau), j =_H i\}$. We have that:

$$A_s = \bigcup_{i=1}^{n} A_s^i$$

Being in this case, s_1 and s_2 are the states where transition of template θ is enabled and picks a channel to send the slice to. We can observe that: $|A_{s_1}^i| = 1$ and $|A_{s_2}^i| = n_i$, for any $i \in [1, k]$. Intuitively: by definition, in the small system we have only one possible concrete transition that picks a server from channel i, while in the big system we have n_i concrete transitions each choosing a different server from channel i. Thus, we can write that: $A_{s_1}^i = \{(s_1, p, \alpha, q)\}$, for some $p \in [0,1]$ and $q \in S_1$, and $A_{s_2}^i = \{(s_1, p_1, \alpha, q_1), \ldots, (s_1, p_{n_i}, \alpha, q_{n_i})\}$, for some $p_1, \ldots, p_{n_i} \in [0,1]$ s.t. $\sum_{j=1}^{n_i} p_j = p$, and $q_1, \ldots, q_{n_i} \in S_2$.

Finally, we can show that, for all i:

$$\begin{aligned} Pr(s_1, \alpha, X) &= Pr(s_1, \alpha, q) = p \\ &= \textstyle\sum_{j=1}^{n_i} Pr(s_2, \alpha, q_j) \\ &= \textstyle\sum_{j=1}^{n_i} p_j \qquad\qquad = Pr(s_2, \alpha, X) \end{aligned}$$

In fact: $\sum_{j=1}^{n_i} p_j = \sum_j g_i(j) \cdot f(i) = f(i) \cdot \sum_j g_i(j) = f(i)$, since $1 = \sum_j g_i(j)$, by our assumptions. ∎

Theorem 3. *Assume MDPs* $\mathcal{M}_1 = \textsc{Client} \parallel \textsc{ProviderAtt}$ *and* $\mathcal{M}_2 = \textsc{Client}' \parallel \textsc{ProviderAtt}$. *Let* L_1 *and* L_2 *be the respective labeling functions, and assume they are invariant with respect to* V. *Then:* $\mathcal{M}_1 \approx \mathcal{M}_2$.

Proof: First of all, let us observe that variables ctr_c^i is always compared with parameter c, the storage capacity. Since by assumption $c > n$, for all $n \in \mathbb{N}_{>0}$, this guard can always be dropped from CLIENT as it is a tautology. What remains is a MDP that updates the variables ctr_c^i but never reads it.

The MDP CLIENT$'$ is identical to CLIENT, except that it does not update the variables ctr_c^i. It is immediate to see that every transition enabled in CLIENT must be enabled also in CLIENT$'$, and viceversa.

Name S_1 the set of states of $\mathcal{M}_1$, and S_2 the set of states of $\mathcal{M}_2$, and name $S = S_1 \uplus S_2$ their disjoint union. Let us define the relation $R \subseteq S \times S$ as follows: $R = \{(s_1, s_2) : s_1 \equiv_V s_2\}$.

By our assumptions, L_1 and L_2 are invariant w.r.t. variables in V, meaning that $s_1 \equiv_V s_2 \Rightarrow L_1(s_1) = L_2(s_2)$. Combined with definition of R we have that $R(s_1, s_2) \Rightarrow L_1(s_1) = L_2(s_2)$.

Since every transition enabled in $\mathcal{M}_1$ is also enabled in $\mathcal{M}_2$, and since they have the same probability, the second requirement of a probabilistic bisimulation holds (see Definition 2).

∎